

Using Qualitative Comparative Analysis to Study Causal Complexity

Charles C. Ragin

In case-oriented research, investigators study the ways in which the different aspects of cases fit together within each case, and they make sense of each case separately (Ragin, earlier in this issue). Although this approach is rich in detail, it highlights individual cases at the expense of knowledge about cross-case patterns. This strategy differs from variable-oriented research, where the key concern is for patterns observable across cases, not the specificity of individual cases. In this article, I bring these two kinds of concerns together. The key to bridging them is a “configurational” view of cases. In this view cases are understood in terms of the aspects they combine, as different configurations of set memberships. Rather than viewing cross-case patterns through the lens of “relationships between variables,” the researcher compares and contrasts configurations. In my presentation of this approach, I emphasize the study of the different combinations of causal conditions that are sufficient for a selected outcome. However, the configurational approach is not limited to the study of causation or causal complexity. It is relevant to any investigation where the number of cases is small enough to permit some degree of familiarity with each case, yet large enough to warrant an interest in cross-case patterns.

First, I explain my approach to causal complexity. This discussion builds on the contrast between studies that focus exclusively on cases displaying a specific outcome and a search for common antecedent conditions, on the one hand, and studies that allow for the possibility that the same outcome can follow from different combinations of conditions, on the other. Next, I present *Qualitative Comparative Analysis* (QCA) (Drass and Ragin 1992), an analytic technique designed specifically for the study of cases as configurations of aspects, conceived as combinations of set memberships. I illustrate QCA

Charles C. Ragin, Ph.D. is Professor of Sociology and of Political Science, Northwestern University, and Professor II of Sociology, University of Oslo, Norway. Address correspondence to Charles C. Ragin, Ph.D., Dept. of Sociology, Northwestern University, Evanston, IL 60208-1330. This article, submitted to *Health Services Research* on February 5, 1999, was revised and accepted for publication on June 17, 1999.

with an analysis of hypothetical data on health maintenance organizations (HMOs). In this analysis, the problem is to identify the different combinations of conditions linked to high rates of staff turnover. Finally, I sketch general issues in the use of QCA.

CAUSAL COMPLEXITY AND CASE-ORIENTED RESEARCH

A common strategy in comparative case-oriented research is to study a small to moderate number of cases in which a specific outcome has occurred. Usually, this design involves a search for antecedent conditions shared by all (or virtually all) instances of the outcome, with an eye to understanding how these conditions fit together to produce the outcome. Cross-case commonalities identified by the investigator provide the basis for constructing a general account of how the outcome comes about. For example, a researcher might study several HMOs that experienced a very high turnover of physicians due to defection to other forms of practice. Common antecedent conditions exhibited by these HMOs, which might include "speed-up" in the patient flow, increased management oversight of referrals to medical specialists, and so on, would contribute to a general understanding of the forces generating the outcome.

Although simple and straightforward, this case-oriented research design is far from problem-free. The most obvious problem is that the investigator's confidence in the causal conditions that he or she has identified increases as the number of instances of the outcome increases. The greater the number of cases examined (e.g., instances of HMOs with defection-related turnover), the more impressive the fact that they share common antecedent conditions (e.g., speed-up, increased oversight, and so on). But as the number of cases increases, so does the difficulty of knowing cases well, making it impossible to become familiar enough with each case to make sound judgments about causally relevant features. Besides, as the number of cases increases, the likelihood that they will share causally relevant features declines. "More cases" almost always means "more heterogeneity."

Another problem with this design is the fact that it is useful only for identifying *necessary* conditions (Dion 1998). When selecting on instances of an outcome, it is not possible to assess the *sufficiency* of causal conditions, at least not in any direct or systematic manner. To assess sufficiency, the researcher must select on instances of the causal condition (or on instances of the relevant

combination of causal conditions), not on instances of the outcome. A cause is sufficient if it is invariably (or almost invariably) followed by the outcome; it is necessary if it is present in all instances of the outcome. (See Ragin in this issue; other common problems with the study of the antecedent conditions shared by instances of an outcome are addressed in Collier 1995; Collier and Mahoney 1996.)

These common problems are not the primary concern of this article. Rather, the focus is on the problem of *causal complexity*: the fact that many of the outcomes that interest social scientists often result from several different, non-overlapping combinations of conditions. For example, several different combinations of conditions may spark "defection-related turnover" in HMOs, and no single antecedent condition may be common to all or even to most instances of the outcome. When causation is this complex—probably the rule and not the exception in the study of social phenomena—then no single causal condition is either necessary or sufficient.

When they face causation this complex, case-oriented researchers are likely to typologize their cases. That is, rather than seeing their cases as instances of the "same thing," they differentiate types of outcomes, for example, types of "defection-related turnover." The analysis of similarities and differences among cases will play a key part in the process of developing types. The investigator will try to identify causal commonalities *within each type* and causal differences *between types*. In effect, the search for commonalities will narrow as different combinations of causal conditions are linked to specific types of the outcome. In the end, a one-to-one correspondence appears between the different types of the outcome and the sets of conditions linked to each type.

There is nothing wrong with resorting to typologies. Often, it is a very creative process that leads to important theoretical insights and advances. However, typologizing may not be the best way, and it is certainly not the only way, to address causal heterogeneity. The danger is that types can proliferate to the point where the number of cases conforming to each type may be quite small. This potential proliferation of types stems in part from an overemphasis on the search for commonalities, understood as necessary conditions; but it is also an effect of the rigidity of the notion underlying it: that each outcome must result from a single causal condition or a single combination of conditions. Once it is acknowledged that the same outcome may result from several different combinations of conditions, then the decision to typologize seems less straightforward.

How should case-oriented researchers address causal heterogeneity?

What should they do when a given outcome results from several different combinations of conditions? QCA is designed specifically for the study of this type of complexity. At its core, it is nonstatistical, but probabilistic criteria can be incorporated at various points in the procedure (see Ragin 1999, 2000). QCA is intended to work hand-in-hand with case-oriented inquiry. It is based on case-oriented knowledge and at the same time provides the means for structuring case-oriented inquiry and for systematizing the results of this type of inquiry. Before describing QCA, I first address the issue of causal complexity, sketching the core of the problem. Even though causal complexity does undermine the examination of single causes (and the use of additive models more generally), I show that causal complexity is not resistant to analysis, *per se*.

NECESSITY, SUFFICIENCY, AND CAUSAL COMPLEXITY

In many fields of social science, the assessment of necessary conditions is very important. Necessary conditions have very powerful policy implications. For example, in some situations it is possible to prevent a bad outcome by removing or blocking one of its necessary conditions. However, it is important to recognize that designs that work backward from multiple instances of an outcome to identify shared antecedent conditions (1) cannot address sufficiency and thus may fail to identify *decisive* causes, and (2) make restrictive assumptions about causation, namely, that theoretically relevant, nontrivial necessary conditions *exist*. I contend that the search for necessary conditions, while often useful, sometimes constricts case-oriented research in unproductive ways. It is much more fruitful to allow for the possibility that a given outcome may follow from a variety of different combinations of theoretically relevant causal conditions.

As an alternative to searching for necessary-but-not-sufficient causal conditions (i.e., studying causal commonalities shared by instances of an outcome), researchers should work forward from causal conditions, especially combinations of conditions, and assess their sufficiency. To assess the sufficiency of a cause or causal combination, the researcher must determine whether or not the cause or combination of causes in question always, or virtually always, produces the outcome in question. Evidence that there are instances in which the cause or causal combination is *not* followed by the outcome challenges the researcher's claim that the cause or causal combination

is sufficient. The assessment of sufficiency, therefore, involves searching for cases that are similar with respect to the relevant causes and then assessing whether or not these cases display the same outcome.

For illustration, suppose two combinations of conditions exist that produce “defection-related turnover” in HMOs: (1) a change in management combined with a speed-up in the patient flow and (2) extensive use of outside medical specialists combined with management appropriation of the power to veto referrals to specialists. Because there are two sufficient combinations of conditions and no single condition is shared by the two combinations, the analysis of necessity would show that no single causal condition (and no causal combination) is necessary. Likewise, the analysis of sufficiency would show that no single causal condition is sufficient. Thus, no single cause is either necessary or sufficient. However, the analysis of the sufficiency of *combinations* of conditions would show that the two listed combinations are sufficient.

The relevant pattern of results is presented in Table 1. In the analysis of either of the two sufficient causal combinations, there would be cases in cell 1 because there is plural causation—more than one way to generate the outcome. However, both of the listed combinations would pass the test of sufficiency. In each test, cell 4 would be void or virtually void of cases, while cell 2 would contain instances of defection-related turnover explained by the causal combination in question. As long as there are cases in cell 2 and no cases (or virtually no cases) in cell 4, then the researcher may argue that the causal combination is sufficient, assuming that theoretical and substantive knowledge support this interpretation. Cells 1 and 3 are not directly relevant to the assessment of sufficiency.

More generally, Table 1 demonstrates that it is possible to assess the sufficiency of causal combinations one at a time, in isolation from one another. This conclusion is important because of its implications for the study of causal complexity. If, as I have argued, we live in a world of great causal complexity, then a common pattern will be for outcomes to result from

Table 1: The Sufficiency of Causal Combinations

	Causal Combination Absent	Causal Combination Present
Outcome Present	1. Cases explained by second causal combination	2. Cases explained by first causal combination
Outcome Absent	3. Not directly relevant to the assessment of sufficiency	4. No cases (or virtually no cases) here

different combinations of causal conditions. Although it might seem that causation this complex would befuddle analytic social science, it is clear from the simple example just presented that the analysis of the sufficiency of causal combinations can proceed in a straightforward manner.

USING "TRUTH TABLES" TO ADDRESS CAUSAL COMPLEXITY

In order to select cases that display the same combination of causal conditions and then to test the sufficiency of that combination, it is necessary first to have a good sense of which cases and which causes are relevant. Sometimes a research literature is especially well developed, and the selection of cases, specification of outcomes, and identification of causally relevant conditions are all relatively unproblematic. In other situations, however, the researcher can formulate a worthwhile selection of cases, outcomes, and causal conditions only through an in-depth analysis of cases. Sometimes it is necessary to constitute relevant cases and their key aspects through a systematic dialogue of ideas and evidence. Researchers progressively refine their understanding of relevant cases as they sharpen the concepts appropriate for studying them (see Ragin 1992, 1994a, 1997, 2000).

In QCA, once the relevant cases, outcomes, and causally relevant conditions have been identified, at least provisionally, the researcher constructs a "truth table" listing the different logically possible combinations of causal conditions along with the cases conforming to each combination (Ragin 1987). One can also view this table as a "property space" (Lazarsfeld 1937; Barton 1955) that locates cases according to their different combinations of causally relevant features. In most applications of QCA, attributes are represented using presence/absence dichotomies. Multichotomies (e.g., race/ethnicity at the individual level) are represented with sets of dichotomies, which can be arranged in a variety of ways depending on the interests of the investigator.

Consider, for example, Table 2, which lists the logically possible combinations of the four causal conditions, mentioned previously, that are relevant to defection-related turnover in HMOs. The four attributes that define this property space are (1) whether or not a recent change has taken place in the ownership and management of the HMO, (2) whether or not management has recently appropriated veto power over all referrals to medical specialists, (3) whether the HMO uses mostly its own medical specialists or relies on specialists outside the HMO, and (4) whether or not the physicians have

Table 2: Hypothetical Truth Table for Causes of “Defection-related Turnover” in HMOs

	<i>Management Change</i>	<i>Management Veto</i>	<i>Outside Specialists</i>	<i>Speed-up</i>	<i>Turnover</i>	N
1	~change	~veto	~outside	~speed-up	no	6
2	~change	~veto	~outside	speed-up	no	8
3	~change	~veto	outside	~speed-up	no	4
4	~change	~veto	outside	speed-up	no	9
5	~change	veto	~outside	~speed-up	?	0
6	~change	veto	~outside	speed-up	?	0
7	~change	veto	outside	~speed-up	yes	4
8	~change	veto	outside	speed-up	yes	7
9	change	~veto	~outside	~speed-up	no	6
10	change	~veto	~outside	speed-up	yes	8
11	change	~veto	outside	~speed-up	?	0
12	change	~veto	outside	speed-up	yes	10
13	change	veto	~outside	~speed-up	no	12
14	change	veto	~outside	speed-up	yes	5
15	change	veto	outside	~speed-up	yes	11
16	change	veto	outside	speed-up	yes	7

Notes: The “~” symbol preceding an attribute name indicates “not” or negation. *Management Change*: whether the HMO has recently experienced a change in ownership and management (change) or not (~change). *Management Veto*: whether the HMO’s management has recently acquired veto power over all referrals to specialists (veto) or not (~veto). *Outside Specialists*: whether the HMO uses outside or intramural (~outside) medical specialists. *Speed-up*: whether the medical staff has had its pace of patient scheduling increased (speed-up) or not (~speed-up).

been subjected to a recent speed-up in the patient flow (the scheduling of more patients per hour). For notational convenience in the discussion that follows, the presence of a condition is denoted by its name; the absence of the condition (negation) is denoted with the “~” (tilde) symbol preceding the attribute name. Four presence/absence dichotomies yield 2^4 (16) logically possible combinations of conditions, each of which is listed in the truth table. The example in Table 2 is very simple. A truth table with six causally relevant conditions, represented with dichotomies, would have 2^6 (64) rows; with eight conditions, there would be 2^8 (256) rows, and so on.

By examining the cases that conform to each combination of causal conditions, represented as a row of the truth table, it is possible for the investigator to evaluate whether or not an adequate set of conditions has been identified. For each row the researcher asks: Do these cases go together? Are they comparable instances, in the context of the present study? For example, the first row, ~change*~veto*~outside*~speed-up (asterisks indicate combinations of characteristics), brings together six HMOs. Viewing these six cases together, the researcher asks whether or not it is reasonable to group them as similar cases in a study of defection-related turnover. If not, then additional attributes should be added to the list of relevant causal conditions, or perhaps the researcher should substitute new attributes for some of the listed attributes. Cases conforming to each row of the table should be evaluated in this manner.

When researchers view their evidence in terms of logically possible combinations of conditions and the cases conforming to each combination, as in Table 2, they also evaluate the cases in each row to see if they display the same outcome, or at least roughly comparable outcomes. For example, a researcher might ask: Are the six cases in the first row similar with respect to whether or not they have experienced defection-related turnover? Each row is evaluated in this manner, so that the researcher can gain some confidence that a viable specification of causal conditions has been realized. Obviously, if the cases in a row display widely divergent outcomes or if they are evenly split between contrasting outcomes, the researcher will examine these cases closely and reformulate his or her specification of causal conditions accordingly. The researcher thus conducts an elaborate dialogue of ideas and evidence that leads to a progressive refinement of understanding of the relevant cases and to a more nuanced elaboration of the relevant causal conditions.

It is important to understand that in QCA the fundamental unit of analysis is the truth table row, along with the cases conforming to each row. Thus, Table 2 should *not* be viewed as a presentation of four presence/absence dichotomies, but rather as a specification of 16 qualitatively distinct causal combinations. More precisely, these 16 combinations represent all of the logically possible *selections* using the four causal conditions. Recall that to assess the sufficiency of a combination of causal conditions, the researcher selects cases with a given combination of conditions and then evaluates whether or not these cases display the same, or roughly the same, outcome. If they all (or virtually all) display the outcome in question, then the evidence supports the argument that the combination of causal conditions in question is sufficient for the outcome. A truth table, in this light, can be seen as an attempt to implement an exhaustive examination of sufficiency. In effect, each row

of the truth table constitutes a different logically possible selection, and each row is evaluated with respect to the outcome. Do the cases in each row agree on the outcome? In short, the truth table provides multiple selections on the independent variables and multiple “tests” of sufficiency.

As I show elsewhere (Ragin 1999, 2000), it is possible to use probabilistic criteria when conducting these tests. This tactic partially ameliorates the problem of contradictory evidence and thus allows for some discordance in outcomes for the cases in a given row. For example, a researcher might test whether or not significantly greater than 80 percent of the cases in a given row exhibit the outcome in question. If they do, then the combination of conditions represented by the row can be described as “almost always sufficient” for the outcome.

ANALYZING CAUSAL COMPLEXITY

Social research usually begins with the goal of explaining some outcome. For example, a researcher might ask why some HMOs in a given set experience defection-related turnover and others do not. Table 2 indicates that HMOs in rows 7, 8, 10, 12, 14, 15, and 16 offer strong evidence of turnover, while those in the other rows do not. How should the researcher describe the key differences between these two sets of HMOs? In other words, what combinations of causal conditions are linked to defection-related turnover?

As presented, the truth table is ready for “logical minimization” because in this example no causal combinations (rows of the truth table) embrace cases with contradictory outcomes (i.e., cases both with and without defection-related turnover in the same row). The goal of the logical minimization of a truth table is to represent—in a logically shorthand manner—the information in the truth table regarding the different combinations of conditions that produce a specific outcome. The procedure described here is identical to that developed by electrical engineers for the minimization of switching circuits (e.g., see Mendelson 1970; Roth 1975). In Table 2, the goal is to specify, in a logically minimal way, the different combinations of change, ~change, veto, ~veto, outside, ~outside, speed-up, and ~speed-up that produce defection-related turnover.

The first step in the minimization process is to select the rows displaying the outcome and compare them with each other. Here the objective is to simplify them through a bottom-up process of paired comparison. These paired comparisons follow a simple rule that mimics experimental design:

combine rows that differ on only one causal condition but produce the same outcome. The 15th and 16th rows, for example, differ only on the fourth condition (speed-up versus ~speed-up), and both produce defection-related turnover. Thus, they can be combined to produce a single, simpler expression. This simpler expression states that if management change, management veto power, and use of outside specialists are all present, the presence/absence of speed-up is irrelevant; defection-related turnover still occurs. This bottom-up procedure continues until no further pair-wise simplifications are possible. For example, the results of the pairing of the 15th and 16th rows (change*veto*outside) can be paired with the results of the pairing of the 7th and 8th rows (~change*veto*outside) to produce the simpler combination, veto*outside. (Asterisks indicate combinations of causal conditions.)

The process of paired comparisons culminates in the production of "prime implicants." In this example two prime implicants result: change*speed-up and veto*outside. Often many more prime implicants are produced than are needed to embrace or cover all of the causal combinations for a particular outcome, and the researcher constructs a "prime implicant chart" (Mendelson 1970; Roth 1974; Ragin 1987) showing the correspondence between the prime implicants (derived from paired comparisons) and the original causal combinations for the outcome of interest drawn from the truth table (i.e., all of the rows with defection-related turnover in Table 2). It is apparent from simple inspection of these results, however, that no prime implicant chart is needed in this example. The two prime implicants derived from paired comparisons are both needed to cover the seven causal combinations for the presence of turnover. Use of the prime implicant chart, if needed, is the final phase of minimization and culminates in a logical equation for the outcome of interest.

The final, reduced equation for the presence of turnover in Table 2 is:

$$\text{turnover} = \text{change*speed-up} + \text{veto*outside}$$

(Addition indicates alternate combinations of conditions, logical *or*.) The equation states simply that there are two combinations of causal conditions that result in defection-related turnover: (1) a change in ownership and management combined with a speed-up of the patient flow, *or* (2) management appropriation of the power to veto all referrals to medical specialists combined with the use of outside specialists. Of course, these hypothetical results are intended only to give the flavor of this type of analysis. The use of empirical

evidence would entail a much more detailed truth table and much greater complexity in the resulting equation.

When only a few causal conditions are examined, QCA can be implemented without the aid of a computer. However, when the number of causal conditions is greater than four, the number of logically possible combinations of conditions increases greatly. For example, an analysis with ten dichotomous conditions results in a truth table with 2^{10} (1,024) rows. Drass and Ragin (1992) offer a microcomputer package called QCA suitable for such analyses.

FURTHER ISSUES IN QUALITATIVE COMPARATIVE ANALYSIS

The demonstration of QCA just offered is based on hypothetical data and is unrealistically simple and straightforward. Several common issues have not been addressed due to space limitations. However, these issues, which will be described briefly here, are discussed in detail elsewhere (Ragin 1987, 1994b, 1995, 1999, 2000). The four main issues are (1) limited diversity, (2) contradictory outcomes, (3) use of categorical data, and (4) analyses that are not explicitly causal in nature.

Limited Diversity

Naturally occurring social phenomena are limited in their diversity. Whenever researchers construct truth tables with a moderate number of conditions, they are likely to find that many of the rows lack cases. How should researchers code the outcome for these rows? There are many ways to answer this question, and different answers can lead to different solutions of the same truth table. It is best to see the rows without cases as "hypothetical combinations" offering a pool of potential "simplifying assumptions" that may be used by researchers to produce more parsimonious solutions of truth tables. However, the incorporation of any simplifying assumption must be justified by the investigator using his or her theoretical and substantive knowledge (Ragin 2000).

The problem of limited diversity also exists in conventional quantitative analysis. In almost all quantitative analyses many regions of the vector space formed by the independent variables are devoid, or virtually devoid, of cases. However, these voids are not addressed in most quantitative analyses, and they do not appear as a problem to researchers, especially those using linear-additive models, the most common type. In effect, conventional quantitative

researchers employ implicit assumptions to deal with these voids. For example, the assumptions of causal additivity and linearity permit researchers to extrapolate predictions, even to regions of the vector space that lack cases. These assumptions remain more or less invisible to both researchers and their audiences.

Contradictory Outcomes

It is very difficult to construct truth tables in which all of the cases in every row agree on the outcome. The challenge of constructing a truth table with a minimum of contradictions often motivates researchers to learn more about their cases and to carry on an enriched dialogue of theory and evidence. Still, contradictory outcomes often remain, even at the end of a very long dialogue. A great deal of the randomness in social life (and social research) is simply outside the scope of any truth table. Individual cases may display (or not display) an outcome for reasons that are entirely specific to the case in question and beyond the purview of any theory. Most applications of QCA employ some sort of rule-of-thumb for resolving contradictions, based on the nature and strength of the evidence and the clarity of the guidance offered by relevant theoretical perspectives. As noted previously in this article, it is possible to use probabilistic criteria to evaluate sufficiency—to guide decisions about rows with contradictory outcomes. For example, a researcher might argue that if significantly greater than 80 percent of the cases in a given row display the outcome (using an alpha of .05), then the combination of conditions specified in the row is “almost always sufficient” for the outcome (see Ragin 1999, 2000).

Categorical Data

As presented here, QCA is limited to categorical data. Yet many of the conditions that interest social researchers are difficult to represent in this way. Consider, for example, one of the causal conditions specified in the hypothetical analysis just described: “speed-up.” How much of an increase in the patient flow is required for an HMO to be placed in the set of HMOs experiencing speed-up? While researchers (and medical personnel) certainly will agree on extreme cases, and maybe they will agree on the classification of most empirical cases, there is still an element of arbitrariness in whatever cut-off value is used to distinguish cases with speed-up from those lacking it. The resolution of this problem can be found in the use of fuzzy sets. For example, the researcher could set up various criteria for membership

in the fuzzy set of HMOs experiencing speed-up and assess their *degree* of membership in this set. Fuzzy membership scores range from 0 (fully out of the set) to 1 (fully in the set), with .5 as the “cross-over point” separating cases that are “more in” versus “more out” of the set in question (see Zadeh 1965, 1972; Ragin 2000). This formulation clearly solves the problem of arbitrary cut-offs. Unfortunately, the current version of QCA (Drass and Ragin 1992) is not equipped to handle fuzzy sets. However, the version that is currently under development will be able to analyze truth tables (conceived as property spaces) composed of fuzzy sets.

Noncausal Analyses

The discussion and the examples in this article all emphasize the analysis of causal complexity. However, the techniques presented are useful for any analysis that focuses on a small to moderate number of cases in a configurational manner. For example, rather than asking, “What causal conditions are linked to defection-related turnover?” a researcher might ask questions about what features go together. The emphasis here would be on features that go together as configurations, not features that correlate. For example, the question “What kinds of HMOs emphasize prevention?” need not be causal in conception. An emphasis on prevention could exist as an important feature within several different “packages” of HMO characteristics. The researcher could examine each package and interpret the place of the emphasis on prevention within each combination of characteristics. Thus, the analysis would have a configurational character, as would the analysis of combinations of causal conditions, but no use would be made of the language of causation or causal reasoning.

SUMMARY: USING QUALITATIVE COMPARATIVE METHODS

The application of QCA to cross-case evidence is carried out in three distinct phases: (1) selecting cases and constructing a truth table that defines their causally relevant characteristics, (2) testing the sufficiency of causal conditions, and (3) deriving and interpreting the results. The summary equations that result from the application of QCA should be viewed as part of the larger dialogue of ideas and evidence (Ragin 1987). The real test of any representation of evidence is how well it helps the researcher and his or her audiences understand specific cases or sets of cases. Broad representations of cross-case

patterns provide maps that guide and facilitate in-depth investigation; they are not substitutes for this type of investigation. Thus, QCA has an implicit fourth phase involving the application of the results to specific cases, but this phase is not part of QCA proper.

In many respects the first phase of QCA is the most difficult. The specification of causally relevant aspects of cases must be clarified and refined to see if the resulting truth table sorts cases in a way that makes sense. At the same time, the researcher must study the cases initially chosen for investigation and evaluate whether or not the set as a whole has integrity. Dropping or adding cases may help the researcher refine his or her specification of conditions while at the same time it may increase the comparability of the cases in the study. Simultaneously, the researcher also examines cases conforming to each row of the truth table with respect to the outcome under investigation, with an eye to their concordance. If cases differ too greatly on the outcome, then either the causal conditions must be re-specified, the set of cases reconstituted, or both.

Once the researcher successfully stabilizes the relevant cases and the truth table, then the assessment of causal sufficiency can proceed. In this phase, the key issue is the definition of sufficiency: How should the test be structured? The answer to this question is shaped in large part by the nature of the evidence and the criteria that are most important to the investigator. Still, in most analyses, it is probably best to work with several definitions of sufficiency and to conduct tests that favor competing criteria. Once these tests are complete, algorithms implemented in the program QCA can be used to analyze and simplify the patterns.

The advance of social scientific knowledge is best served when scholars make as few assumptions about causation as possible, especially at the outset of an investigation. When scholars assume maximum causal complexity—that different combinations of causes may produce the same outcome—they assume that no single cause is either necessary or sufficient. As I have shown, analytic social science is possible even when causal complexity is great. The analysis of causal complexity, in turn, is greatly facilitated by the use of QCA. This approach offers a powerful way to assess the sufficiency of causal conditions, a task that is outside the domain of conventional quantitative analysis (see Ragin, "The Distinctiveness of Case-Oriented Research," in this issue).

As noted previously, necessary conditions have very powerful policy implications, especially when the goal is to prevent some unwanted outcome. Remove one of the necessary conditions for an outcome and the outcome is blocked. Thus, an emphasis on necessary conditions highlights the fragility

of social action and the ease of disruption. By contrast, the study of sufficiency, and especially of the alternative combinations of conditions sufficient for an outcome, emphasizes the open-ended nature of social action. When researchers show that many different ways exist to reach the same outcome, they show, in effect, the creative and inventive side of social action. Far from emphasizing possibilities for disruption, they unveil different pathways.

REFERENCES

- Barton, A. H. 1955. "The Concept of Property Space in Social Research." In *The Language of Social Research*, edited by P. F. Lazarsfeld and M. Rosenberg. Glencoe, IL: Free Press.
- Collier, D. 1995. "Translating Quantitative Methods for Qualitative Researchers: The Case of Selection Bias." *American Political Science Review* 89 (2): 461–66.
- Collier, D., and J. Mahoney. 1996. "Insights and Pitfalls: Selection Bias in Qualitative Research." *World Politics* 49 (1): 56–91.
- Dion, D. 1998. "Evidence and Inference in the Comparative Case Study." *Comparative Politics* 30 (1): 127–45.
- Drass, K., and C. C. Ragin. 1992. *QCA: Qualitative Comparative Analysis*. Evanston, IL: Institute for Policy Research, Northwestern University.
- Lazarsfeld, P. F. 1937. "Some Remarks on Typological Procedures in Social Research." *Zeitschrift Fur Sozialforschung* 6 (1): 119–39.
- Mendelson, E. 1970. *Boolean Algebra and Switching Circuits*. New York: McGraw-Hill.
- Ragin, C. C. 2000. *Fuzzy-Set Social Science*. Chicago: University of Chicago Press.
- . 1999. "The Logic of Qualitative Comparative Analysis." *International Review of Social History* forthcoming.
- . 1997. "Turning the Tables: How Case-Oriented Research Challenges Variable-Oriented Research." *Comparative Social Research* 16 (1): 27–42.
- . 1995. "Using Qualitative Comparative Analysis to Study Configurations." In *Computer-Aided Qualitative Data Analysis*, edited by U. Kelle, pp. 177–89. London: Sage.
- . 1994a. *Constructing Social Research: The Unity and Diversity of Method*. Thousand Oaks, CA: Pine Forge Press.
- . 1994b. "Introduction to Qualitative Comparative Analysis." In *The Comparative Political Economy of the Welfare State*, edited by T. Janoski and A. Hicks, pp. 299–319. New York: Cambridge University Press.
- . 1992. "Casing and the Process of Social Research." In *What Is a Case? Exploring the Foundations of Social Inquiry*, edited by C. C. Ragin and H. S. Becker, pp. 217–26. New York: Cambridge University Press.
- . 1987. *The Comparative Method: Moving Beyond Qualitative and Quantitative Strategies*. Berkeley, CA: University of California Press.
- Roth, C. 1975. *Fundamentals of Logic Design*. St. Paul, MN: West.
- Zadeh, Lofti A. 1972. "A Fuzzy-Set-Theoretic Interpretation of Linguistic Hedges." *Journal of Cybernetics* 2 (1): 4–34.
- . 1965. "Fuzzy Sets." *Information Control* 8 (3): 338–53.